

SpaceFibre Interface and Routing Switch IP Cores

SpaceFibre, Long Paper

Albert Ferrer Florit, Alberto Gonzalez Villafranca

STAR-Barcelona S.L.

Sant Cugat, Spain

albert.ferrer@star-dundee.com

Steve Parkes, Chris McClements

STAR-Dundee Ltd.

Dundee, UK

Abstract—SpaceFibre is a technology specifically designed for use onboard spacecraft that provides point to point and networked interconnections at Gigabit rates with Quality of Service (QoS) and Fault Detection, Isolation and Recovery (FDIR). SpaceFibre is backwards compatible with SpaceWire, allowing existing SpaceWire equipment to be incorporated into a SpaceFibre network without modifications at packet level.

In this work we present the family of SpaceFibre IP Cores developed by STAR-Dundee. It is composed of three different IPs: the Single-Lane Interface, the Multi-Lane Interface and the Routing Switch. The IP Cores are fully compliant with the SpaceFibre standard and have been carefully implemented to optimise their performance and minimise their footprint on radiation-tolerant FPGAs (e.g. RTAX, RTG4, BRAVE or Virtex-5QV) and ASICs. They have also been validated on commercial FPGAs (e.g. Iglou2, Spartan, Virtex, Kintex, etc.).

The Single-Lane Interface IP offers in a compact design (~3% of the RTG4/Virtex-5QV) the maximum possible line rates provided by embedded or external transceivers (i.e. 3.125 Gbps in RTG4, 4.25 Gbps in Virtex-5QV and 2.5 Gbps in RTAX using the TLK2711-SP transceiver). The Multi-Lane Interface IP allows much higher data rates and adds all the advantages of combining multiple lanes without multiplying the resources required (e.g. ~5-6% for 3 lanes in RTG4/Virtex-5QV). The SpaceFibre Routing Switch IP Core is a scalable, fully configurable non-blocking router, allowing to select the number of virtual channels and ports. This routing switch implements path and logical addressing, group adaptive routing, virtual networks, time distribution and message broadcast. A router of 4 ports each with 4 virtual channels typically requires less than 20% of an RTG4, including the SpFi interfaces.

The IP Cores presented in this article provide the building blocks for creating the next generation of onboard networks with in-built QoS and FDIR mechanisms, and are currently being implemented in several missions and products all over the world. We analyse the performance and capabilities of the different IP Cores, and discuss the resources required depending on several parameters such as the number lanes, ports, virtual channels and virtual networks.

Index Terms—SpaceFibre, Single-Lane, Multi-Lane, Routing Switch, IP Core, FPGA, Radiation Tolerant, RTG4, Virtex-5QV, Networking, Spacecraft Electronics

I. INTRODUCTION

SpaceFibre (SpFi) [1-2] is a new technology for use onboard spacecraft that provides point-to-point and networked interconnections at Gigabit rates with QoS. SpFi interoperates seamlessly with a SpaceWire (SpW) [3-4] network over virtual channels (VCs), as it uses the same packet definition. It provides broadcast capabilities and is able to operate over a copper or fibre optic communication medium. SpFi will be released as an ECSS standard later this year.

New generation payloads, such as SAR and multi-spectral imaging instruments, require the use of multiple parallel high-speed links to fulfil the increasing bandwidth requirements. To accommodate these needs, SpFi supports multi-lane operation, thus allowing data to be sent over several individual physical lanes to enhance throughput and robustness. The multi-lane operation allows much higher data rates through lane aggregation, supporting any number of lanes (up to 16) and unidirectional operation. This effectively multiplies the throughput of the interface by combining several lanes into a single link. Furthermore, when a lane fails the multi-lane mechanism supports hot redundancy and graceful degradation by automatically spreading traffic over the remaining working lanes. Thus multi-lane provides all the advantages of combining multiple lanes without multiplying the resources required.

The SpFi Network layer is responsible for transferring packets over a SpFi link or network. The information to be sent is packaged in the same format as SpW: <Destination Address> <Cargo> <End of Packet Marker>. It uses the same routing concepts as SpW including both path and logical addressing. The SpFi Network layer defines the concept of Virtual Network (VN). VNs are built from the interconnection between VCs of different ports. These VNs enable the creation of highly flexible SpFi routing switches comprising a number of SpFi interfaces and a fully configurable, non-blocking, high performance, routing switch. This routing switch can theoretically support an arbitrary number of VNs, each effectively behaving like independent SpW networks capable of working at multi-Gbps rates.

This paper describes the SpFi Single-Lane IP Core features and performance in section II. Section III analyses the SpFi Multi-Lane IP. Section IV describes the characteristics of the

SpFi Routing Switch IP. Finally, conclusions are presented in section V.

II. SPACEFIBRE SINGLE-LANE INTERFACE IP CORE

The STAR-Dundee SpaceFibre Single-Lane Interface IP Core is compliant with the SpFi standard with the exception of the Multi-Lane layer, which has been added in a separate IP (presented in Section III).

A. Characteristics

This IP has been designed to minimise as much as possible the designer effort of adding the SpFi interface into a design. It is provided with a protocol agnostic data interface, so that no prior knowledge of the SpFi standard is required. Simple data interfaces based on standard 32-bit input and output FIFO interfaces are used. Specifically, they follow the AXI4-Stream protocol [5], which is an industry standard commonly used. This AXI4-Stream interface allows using independent user-defined read and write clocks. Inside them the clock synchronisation between the user local clock domains and the SpFi IP clock domain is done.

The IP can be configured using generics. Different properties can be configured:

- The type of transceiver interface
- The number of VCs
- The size of the VC buffers
- The type of Management interface

There are different high-speed transceiver interface options that provide simple set of signals to be directly connected to the selected transceiver. However, for user convenience different wrappers are supplied for the different transceiver interfaces. For example, old FPGA technologies require external transceivers. In the IP there is a dedicated TLK2711-SP (Wizardlink) [6] interface with all the required signals readily available at the top level. Regarding newer FPGA technologies, there is a specific interface for the RTG4 SerDes, with the 8B/10B encoding (20-bit interface), clock correction and symbol alignment done inside the IP. In this case, three clocks are required by the IP. If we define the lane speed as L_s then two clocks with frequencies of $L_s/20$ and $L_s/40$ are needed for transmission and general operation (e.g. 156.25 & 78.125 MHz @ 3.125 Gbps). The slowest clock is the main clock of the IP and is used by most of the logic. In addition, a recovered clock working at $L_s/20$ is required in reception.

In the case of Xilinx devices (e.g. Virtex-5QV), a simpler interface (32+4 bits) is available that benefits from the additional capabilities offered by their in-built transceivers. The 8B/10B encoding, clock correction and symbol alignment are done inside the Xilinx transceiver. Only a single clock is then required by the IP, with the exception of the AXI clocks of the VC interfaces. This clock operates at a frequency of $L_s/40$ (e.g. 78.125 MHz @ 3.125 Gbps).

The QoS is independently and dynamically configurable for each VC. Three different mechanisms can be configured and combined: scheduling, priority and bandwidth reservation. The FDIR mechanisms automatically recover from transient and persistent errors on the SpFi link. A transient error takes less

than 2 μ sec to recover and does not affect the user data rate thanks to the embedded buffering inside the IP. Other protections against errors include data and broadcast babbling idiot protection. Tests have been done to validate that data integrity is maintained for Bit Error Rates (BER) better than 10^{-5} . The lane is automatically disconnected when the BER is worse than 10^{-5} to prevent a potential protocol breakdown.

The IP has been optimised for low latency operation. Moreover, it offers an 80-bit interface that makes use of the broadcast message mechanism in SpFi to deliver ultra-low latency short messages. These messages present less than 400 ns of guaranteed latency for a few meters of cable.

A simple management interface allows real time configuration of the control parameters and access to the IP status. Furthermore, this interface also includes optional statistics and debug signals. Two different types of management interfaces can be selected: APB bus or signal bus. Having independent signals for each status and configuration field is useful when implementing the IP in an FPGA design that needs direct access to the IP. The alternative is to access to all these fields over an APB bus. This greatly simplifies the interface for designs that use a CPU or want a centralised access point to several interfaces, for example. The APB bus also comes with an independent clock for convenience, and manages internally the clock synchronisation with the clocks of the SpFi IP.

Finally, power management options have been considered. For example, there is the possibility of start one end of the link in a low-power mode waiting for the other end to become active. A block diagram of the interconnection of the IP Core is shown in Fig. 1.




Fig. 1. SpaceFibre Single-Lane interconnection block diagram

B. Performance

The SpFi Single-Lane IP design has been validated in the main FPGA families relevant to space, such as Microsemi RTG4 [7], and Xilinx Virtex-5QV [8] and RTAX. It offers in a compact design the maximum possible line rates in radiation-tolerant FPGAs provided by embedded or external transceivers (i.e. 3.125 Gbps in RTG4, 4.25 Gbps in Virtex-5QV and 2.5

Gbps in RTAX using the TLK2711-SP device). The IP will also be compatible with the NanoXplore BRAVE Large FPGA [9], and can be implemented in commercial FPGA families such as SmartFusion, Spartan, Kintex, Zynq, etc.

The IP has been extensively tested in the RTG4 and Virtex-5QV to guarantee timing closure for the whole temperature and voltage range (i.e. fast & slow corners) required by the space devices, including the EDAC in the memories and the Single Event Transient (SET) filtering enabled. These radiation mitigation techniques reduce the maximum speed of the designs and can sometimes give problems to meet the targeted speeds. Additionally, the IP has been tested under radiation in an RTG4. The information gathered during the test campaign has allowed for further refining the operation under radiation of the IP and the associated reference design in the RTG4. This IP Core is also ready for ASIC implementation. It has been used in the new RC64 many-core DSP ASIC developed by Ramon Chips. This processor features 12 SpFi ports with line rates of up to 6.25 Gbps [10].

The RTG4/Virtex-5QV resources required by the SpFi interface including transmit and receive FIFOs are detailed in Table I for different numbers of VCs. Note that adding a VC to the design has a small impact in the overall resource usage.

TABLE I. SPFI SINGLE-LANE RESOURCE USAGE

	RTG4			Virtex-5QV		
	LUT	DFP	LSRAM Block	LUT	DFP	RAM Block
1 VC	3944 2.6%	2818 1.9%	4 1.9%	2750 3.4%	2365 2.9%	4 1.3%
2 VCs	4454 2.9%	3197 2.1%	6 2.9%	3138 3.8%	2596 3.2%	6 2.0%

III. SPACEFIBRE MULTI-LANE INTERFACE IP CORE

Multi-lane is an optional capability of a SpFi link defined in the Multi-Lane layer of the SpFi protocol stack. As shown in Fig. 2, the Multi-Lane layer is defined between the Data Link layer and the Lane layer implemented for each available lane.

The Data Link layer provides QoS and flow control for a SpFi link. It frames the information to be sent over the link and also provides error recovery capabilities, detecting any frames or control words that go missing or arrive containing errors and resending them. On the other hand, the Lane Layer establishes a connection across a SpFi lane and ensures that bit, symbol and word synchronizations are achieved and that the two ends of the lane are both ready to send and receive data. The Lane Layer also encodes data and control words into 8B/10B symbols, sends and receives symbols over the lane and decodes the received symbols into data and control words. The Multi-lane layer coordinates the operation of multiple lanes as a single SpFi link, providing higher data throughput and redundancy. Because the logic that initializes a lane and monitors its status is located below the multilane layer, each lane can be initialized and operated independently of each other.




Fig. 2. Multi-Lane layer in the SpaceFibre layer stack

This architecture also supports the implementation of graceful degradation, which means that in the event of one or more lanes failing, traffic is spread over the remaining working lanes automatically. When combined with the Data Link layer QoS, the bandwidth allocated to lower priority VCs is reduced when required to ensure that most important information gets through and deterministic traffic is maintained.

A. Characteristics

The SpFi Multi-Lane IP core has been designed to be easy to use, with minimum configuration signals. On top of the features of the Single-Lane IP Core already mentioned in Section II, the Multi-Lane IP also features the additional capabilities specifically related to the Multi-Lane layer.

The number of lanes is fully configurable, with any number of lanes supported (up to 16). Each lane can independently be selected as uni/bidirectional and hot/cold redundant. Single-Lane SpFi implementations must be bidirectional even if the end-user data flow is unidirectional, because of the feedback required by the protocol. However, in a Multi-Lane implementation only one bidirectional lane is enough for protocol related information. Therefore, other lanes can be unidirectional to save power and mass in asymmetric data flows.

In the 4-lane configuration example of Fig. 3, bidirectional lane 2 can be set to a unidirectional lane for power saving reasons. Unidirectional Lane 4 can be enabled when one lane fails or a higher data rate is required. It is also possible to send data using all four lanes and add an additional one configured as a hot redundant lane, which only sends data when another lane fails.




Fig. 3. Unidirectional and bidirectional lanes interconnection example

The IP Core has been designed to fully support the redundancy capabilities offered by the Multi-Lane. Hot redundant lanes recover not only from transient errors like in the Single-Lane case, but also from lane failures in less than

2 μ s without user intervention and without any data loss. This time is close to the round trip delay of the lane, so short than the data flow of the user is not affected when a lane fails, as the data is internally buffered during the time it takes to resume sending data. In case of lane failure when not using a hot redundant lane, there is an automatic graceful degradation of the link bandwidth, with higher priority VCs being less affected. The QoS mechanism ensures that the most important data is sent first. If a cold redundant lane is available it will be activated in less than 20 μ s, providing warm redundancy.

Bandwidth overprovision and dynamic power management are also possible. These capabilities are very useful for space applications where strict power constrains and a high level of reliability is required on the harsh space environment.

The width of the AXI4-Stream interface for the VCs is configurable in multiples of the SpFi word size (32-bits). This allows supporting slower user clocks and still being able to send or receive data at the maximum speed over a single VC. For example, a 64-bit width can be used to operate a link with 2 lanes at the same clock speed of a single-lane link.



Fig. 4. SpFi Multi-Lane IP interconnection test between RTG4 and Spartan FPGAs

B. Performance

Table II provides the resource usage for two radiation hardened FPGAs, Microsemi RTG4 and Xilinx Virtex-5QV for different number of lanes and VCs. Individual lanes can operate up to

3.125 Gbps, with aggregate rates of up to 6.25 Gbit/s using 2 lanes or up to 12.5 Gbit/s using 4 lanes, for example. This Multi-Lane IP has been tested in the RTG4 and Virtex-5QV to guarantee timing closure with memories using EDAC and the Single Event Transient (SET) filtering enabled in worst case (slowest) scenario. It has also been tested under radiation running in an RTG4, just like has been done with the Single-Lane. Fig. 4 shows the Multi-Lane IP integrated in an RTG4 transferring data to a PXI board with a Xilinx commercial device.

TABLE II. SPFI MULTI-LANE RESOURCE USAGE

	RTG4				Virtex-5QV		
	LUT	DFP	LSRAM Block	μ SRAM Block	LUT	DFP	RAM Block
2 Lanes	6794	4534	8	3	3858	3938	8
1 VC	4.5%	3.0%	3.8%	1.4%	4.7%	4.8%	2.7%
2 Lanes	7736	5285	12	3	4503	4382	12
2 VCs	5.1%	3.5%	5.7%	1.4%	5.5%	5.3%	4.0%
3 Lanes	8771	5593	8	4	4770	4849	8
1 VC	5.8%	3.7%	3.8%	1.9%	5.8%	5.9%	2.7%
3 Lanes	9691	6346	12	4	5416	5226	12
2 VCs	6.4%	4.2%	5.7%	1.9%	6.6%	6.4%	4.0%

IV. SPACEFIBRE ROUTING SWITCH IP CORE

The router architecture is built around a non-blocking routing switch matrix with a configurable number of ports. Each port implements a number of VCs. Each VC has an associated VN number. The switch matrix interconnects one or more VCs with the same VN number, but each of these VCs must be located in a different port.

Packets belonging to different VNs never interfere with one another and do not impact the throughput and latency within the routing switch matrix. On the other hand, when multiple packets in the same VN need to be transferred to the same output port, packet-by-packet, round-robin arbitration is performed, similarly to a SpW router.

A. Characteristics

The STAR-Dundee SpaceFibre Routing Switch Core (SpFi Router for short) IP is a scalable, fully configurable non-blocking router. The IP is very flexible, allowing to select the number of VCs and ports, among other options, using generics. This router implements path and logical addressing, group adaptive routing, VNs, time distribution and message broadcast. In addition, it also fully supports the QoS and FDIR capabilities native to SpFi.

The IP Core has been optimised to work with the existing high-performance radiation-tolerant FPGAs using an arbitrary number of SpFi and internal ports. Each port supports up to eight VCs, each with its own QoS parameters. Furthermore, the Router includes an optional configuration port with RMAP protocol. The maximum number of VNs currently supported is 64, but each of these VNs is completely flexible: any VC of any port can be configured to any VN. The IP Core can also be implemented in ASIC technologies. The high flexibility of the SpFi Router IP Core ensures different user needs can be accommodated with ease.

The Router IP presents a deterministic low latency switching. Round-robin packet arbitration can only occur within each VN. When arbitration is required, it only takes place when two or more VCs request to access to the same output port within the same VN. The Router offers two configurable time-outs. These two timers are required to prevent the blockage of a VN in cases when the source or the sink stall in the middle of a packet, or when there is a babbling node. When timing out the router performs automatic packet spilling of the packet causing the blockage.

The VNs can be configured statically during FPGA programming using VHDL constants, or they can be dynamically modified by the user using logic connected to the configuration port or the RMAP protocol, which can be accessed over one of the ports of the Router.

B. Performance

The Router IP has been validated in commercial FPGAs, e.g. Xilinx Spartan (Fig. 4) or Kintex families, and in space-grade FPGAs such as the RTG4 (Fig. 5). It supports lane rates up to 3.125 Gbit/s in RTG4 or Virtex-5QV, with guaranteed timing closure in RTG4 or Virtex-5QV with EDAC and SET filter enabled and worst-case conditions.




Fig. 4. 9-port SpFi Routers mounted in a cPCI rack forming a SpFi network




Fig. 5. RTG4 PXIe board used to validate the SpFi Router

The actual resources used by the complete SpFi Router design depend on the FPGA and the specific user configuration. In the case of the RTG4, Table III presents the resource usage for two different Router sizes, and for two different types of VN configuration. The values reported already include the SpFi Single-Lane IP Cores used by each of the ports and an additional RMAP configuration port. The *Static VNs* indicated in the table consists of a version in which the VNs are pre-defined during the design phase and hardcoded into the Router design. This can be enough for certain use-cases in which the VNs are not going to change. Note that the values for the *Static VNs* column have been obtained for the most demanding VN configuration. Therefore, the numbers can be further reduced if simpler or fewer VNs are required. On the other hand, if a fully-configurable set of VNs is required, then the values of the *Dynamic VNs* column applies. In this case the VNs can be dynamically modified over RMAP through the configuration port.

TABLE III. SPFI ROUTER RESOURCE USAGE IN THE RTG4 FPGA

	Static Virtual Networks			Dynamic Virtual Networks		
	LUT	DFP	LSRAM Block	LUT	DFP	RAM Block
4 Ports	30990	17735	30	58226	34371	46
4 VCs	20.4%	11.7%	14.4%	38.4%	22.6%	22.0%
9 Ports	64085	35384	80	98339	57099	116
4 VCs	42.2%	23.3%	38.3%	64.8%	37.6%	55.5%

V. CONCLUSIONS

The STAR-Dundee SpaceFibre Single-Lane, Multi-Lane and Routing Switch IP Cores have been designed to be easy to implement in radiation-hardened FPGAs. In this article we have detailed the performance and capabilities of the different IP Cores, and discussed the resources required depending on several parameters such as the number lanes, ports, VCs and VNs.

It is possible to integrate a SpFi Single-Lane interface inside currently available radiation-hardened FPGAs by using only a few percentage (~3%) of the device resources while getting the maximum available speed out of their integrated transceivers. The multi-lane capability dramatically increases the data throughput of SpFi and the addition of multiple lanes provide hot and warm redundancy, or graceful degradation of the link bandwidth when no redundant lanes are available. The SpFi Multi-Lane IP Core provides all these new multi-lane capabilities to RTG4 and Virtex-5QV FPGAs with low resource usage and high performance. Therefore, if more bandwidth or additional robustness is required out of a SpFi link, the Multi-Lane IP is available at the cost of a small increase in the resource usage (e.g. 5-6% for three lanes). Finally, the SpFi Routing Switch can also be integrated in space-grade FPGAs with a size greatly varying depending on the number of VCs and ports, and whether Static or Dynamic VNs configuration is requested.

All the IP Cores have been verified in simulation and subsequently validated in hardware prototypes. Both commercial and the main radiation-hardened FPGAs have been

used for these validation activities, ensuring full compatibility and defining an easy adoption path for this technology.

All the STAR-Dundee family of SpFi IP Cores come with a reference design for RTG4 (Libero), Virtex-5QV (ISE) and Kintex-7 (Vivado) that can directly be implemented in the FPGA to assist the end-user and allow an easy adoption. A comprehensive end-user test bench for ModelSim/Quarta simulators is also provided, which can be used as a reference for test integration. Finally, a Bus Functional Model (BFM) is included to assist in system-level design validation activities.

Together, these IPs provide the building blocks for creating the next generation of onboard networks, and are currently being implemented in FPGA and ASIC designs by different missions and products all over the world.

ACKNOWLEDGMENT

The research reported in this paper has been supported in part by the following organisations and contracts:

- The European Space Agency under ESA contract numbers 4000102641 (SpaceFibre Demonstrator) and 17938/03/NL/LvH (SpaceFibre), and Airbus DS sub-contract number G010003963 (NGMMA).
- The European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement numbers 263148 (SpaceWire-RT) and 284389 (VHiSSI).
- The UK Space Agency and CEOI-ST under University of Leicester contract number RP10G0348A02 (SUNRISE), RP10G0348B206 (SpaceFibre-VPX) and RP10G0348A207 (SpaceDSP).
- The Centre for Earth Observation Instrumentation under University of Leicester contract numbers RP10G0327D13 (LOCUS Elegant Breadboard).

REFERENCES

- [1] S. Parkes, A. Ferrer Florit and A. Gonzalez Villafranca, "SpaceFibre Standard", Draft K2, University of Dundee, December 2017.
- [2] S. Parkes et al, "SpaceFibre: Multi-Gigabit/s Interconnect for Spacecraft On-board Data Handling", IEEE Aerospace Conference, Big Sky, Montana, 2015.
- [3] ECSS Standard ECSS-E-ST-50-12C, "SpaceWire, Links, Nodes, Routers and Networks", Issue 1, European Cooperation for Space Data Standardization, July 2008, available from <http://www.ecss.nl>.
- [4] S. Parkes, "SpaceWire Users Guide", ISBN: 978-0-9573408-0-0, STAR-Dundee, 2012, available from <https://www.star-undee.com/knowledge-base/spacewire-users-guide> (last visited 1st May 2018).
- [5] ARM, "AMBA AXI Protocol Version 2.0 Specification", ARM, Ed., 2010.
- [6] Texas Instruments Inc., "1.6-Gpbs to 2.5-Gbps Class V Transceiver," Texas Instruments Inc., 2018, available from <http://www.ti.com/product/tlk2711-sp> (last visited 1st May 2018).
- [7] Microsemi, "Product Brief RTG4 FPGAs PB0051", Microsemi, 2017, available from https://www.microsemi.com/document-portal/doc_download/134430-pb0051-rtg4-fpgas-product-brief (last visited 1st May 2018).
- [8] Xilinx, "Radiation-Hardened, Space-Grade Virtex-5QV Family Overview DS192". Xilinx, Oct. 11, 2018, available from https://www.xilinx.com/support/documentation/data_sheets/ds192_V5QV_Device_Overview.pdf (last visited 1st May 2018).
- [9] J. Le Mauff, "From eFPGA cores to RHBD SoC FPGA", Space FPGA Users Workshop, 4th edition, ESA/ESTEC, Noordwijk, 2018.
- [10] R. Ginosar, P. Aviely, T. Israeli and H. Meirov, "RC64: High Performance Rad-Hard Manycore", IEEE Aerospace Conference, Big Sky, Montana, 2016.